

METHOD FOR PREDICTING FILE DOWNLOAD TIME FROM MIRRORED DATA CENTERS IN A GLOBAL COMPUTER NETWORK

ABSTRACT OF THE DISCLOSURE

5 An intelligent traffic redirection system performs global load balancing for Web sites located at mirrored data centers. The system relies on a network map that is generated continuously for the user-base of the entire Internet. Instead of probing each local name server (or other host) that is connectable to the mirrored data centers, the network map identifies connectivity with respect to a much smaller set of proxy points,
10 called "core" (or "common") points. A core point then becomes representative of a set of local name servers (or other hosts) that, from a data center's perspective, share the point. Once core points are identified, a systematic methodology is used to estimate predicted actual download times to a given core point from each of the mirrored data centers. Preferably, ICMP (or so-called "ping" packets) are used to measure roundtrip time (RTT) and latency between a data center and a core point. Using such data, an average latency is
15 calculated, preferably using an exponentially time-weighted average of all previous measurements and the new measurement. A similar function is used to calculate average packet loss. Using the results, a score is generated for each path between one of the data centers and the core point, and the score is representative of a file download time. Preferably, the score is generated by modifying an average latency with a penalty factor
20 dependent on the time-weighted average loss function. Whichever data center has the best score (representing the best-performing network connectivity for that time slice) is then associated with the core point.